

# Cross-Language Information Retrieval for Technical Documents

Atsushi Fujii and Tetsuya Ishikawa  
University of Library and Information Science  
1-2 Kasuga Tsukuba 305-8550, JAPAN  
{fujii,ishikawa}@ulis.ac.jp

## Abstract

This paper proposes a Japanese/English cross-language information retrieval (CLIR) system targeting technical documents. Our system first translates a given query containing technical terms into the target language, and then retrieves documents relevant to the translated query. The translation of technical terms is still problematic in that technical terms are often compound words, and thus new terms can be progressively created simply by combining existing base words. In addition, Japanese often represents loanwords based on its phonogram. Consequently, existing dictionaries find it difficult to achieve sufficient coverage. To counter the first problem, we use a compound word translation method, which uses a bilingual dictionary for base words and collocational statistics to resolve translation ambiguity. For the second problem, we propose a transliteration method, which identifies phonetic equivalents in the target language. We also show the effectiveness of our system using a test collection for CLIR.

## 1 Introduction

Cross-language information retrieval (CLIR), where the user presents queries in one language to retrieve documents in another language, has recently been one of the major topics within the information retrieval community. One strong motivation for CLIR is the growing number of documents in various languages accessible via the Internet. Since queries and documents are in different languages, CLIR requires a translation phase along with the usual monolingual retrieval phase. For this purpose, existing CLIR

systems adopt various techniques explored in natural language processing (NLP) research. In brief, bilingual dictionaries, corpora, thesauri and machine translation (MT) systems are used to translate queries or/and documents.

In this paper, we propose a Japanese/English CLIR system for technical documents, focusing on translation of technical terms. Our purpose also includes integration of different components within one framework. Our research is partly motivated by the “NACSIS” test collection for IR systems (Kando et al., 1998)<sup>1</sup>, which consists of Japanese queries and Japanese/English abstracts extracted from technical papers (we will elaborate on the NACSIS collection in Section 4). Using this collection, we investigate the effectiveness of each component as well as the overall performance of the system.

As with MT systems, existing CLIR systems still find it difficult to translate technical terms and proper nouns, which are often unlisted in general dictionaries. Since most CLIR systems target newspaper articles, which are comprised mainly of general words, the problem related to unlisted words has been less explored than other CLIR subtopics (such as resolution of translation ambiguity). However, Pirkola (1998), for example, used a subset of the TREC collection related to health topics, and showed that combination of general and domain specific (i.e., medical) dictionaries improves the CLIR performance obtained with only a general dictionary. This result shows the potential contribution of technical term translation to CLIR. At the same time, note that even domain specific dictionaries

<sup>1</sup><http://www.rd.nacsis.ac.jp/~ntcadm/index-en.html>

do not exhaustively list possible technical terms. We classify problems associated with technical term translation as given below:

- (1) technical terms are often compound word, which can be progressively created simply by combining multiple existing morphemes (“base words”), and therefore it is not entirely satisfactory to exhaustively enumerate newly emerging terms in dictionaries,
- (2) Asian languages often represent loanwords based on their special phonograms (primarily for technical terms and proper nouns), which creates new base words progressively (in the case of Japanese, the phonogram is called *katakana*).

To counter problem (1), we use the compound word translation method we proposed (Fujii and Ishikawa, 1999), which selects appropriate translations based on the probability of occurrence of each combination of base words in the target language. For problem (2), we use “transliteration” (Chen et al., 1998; Knight and Graehl, 1998; Wan and Verspoor, 1998). Chen et al. (1998) and Wan and Verspoor (1998) proposed English-Chinese transliteration methods relying on the property of the Chinese phonetic system, which cannot be directly applied to transliteration between English and Japanese. Knight and Graehl (1998) proposed a Japanese-English transliteration method based on the mapping probability between English and Japanese *katakana* sounds. However, since their method needs large-scale phoneme inventories, we propose a simpler approach using surface mapping between English and *katakana* characters, rather than sounds.

Section 2 overviews our CLIR system, and Section 3 elaborates on the translation module focusing on compound word translation and transliteration. Section 4 then evaluates the effectiveness of our CLIR system by way of the standardized IR evaluation method used in TREC programs.

## 2 System Overview

Before explaining our CLIR system, we classify existing CLIR into three approaches in

terms of the implementation of the translation phase. The first approach translates queries into the document language (Ballesteros and Croft, 1998; Carbonell et al., 1997; Davis and Ogden, 1997; Fujii and Ishikawa, 1999; Hull and Grefenstette, 1996; Kando and Aizawa, 1998; Okumura et al., 1998), while the second approach translates documents into the query language (Gachot et al., 1996; Oard and Hackett, 1997). The third approach transfers both queries and documents into an interlingual representation: bilingual thesaurus classes (Mongar, 1969; Salton, 1970; Sheridan and Ballerini, 1996) and language-independent vector space models (Carbonell et al., 1997; Dumais et al., 1996). We prefer the first approach, the “query translation”, to other approaches because (a) translating all the documents in a given collection is expensive, (b) the use of thesauri requires manual construction or bilingual *comparable* corpora, (c) interlingual vector space models also need comparable corpora, and (d) query translation can easily be combined with existing IR engines and thus the implementation cost is low. At the same time, we concede that other CLIR approaches are worth further exploration.

Figure 1 depicts the overall design of our CLIR system, where most components are the same as those for monolingual IR, excluding “translator”.

First, “tokenizer” processes “documents” in a given collection to produce an inverted file (“surrogates”). Since our system is bidirectional, tokenization differs depending on the target language. In the case where documents are in English, tokenization involves eliminating stopwords and identifying root forms for inflected words, for which we used “WordNet” (Miller et al., 1993). On the other hand, we segment Japanese documents into lexical units using the “ChaSen” morphological analyzer (Matsumoto et al., 1997) and discard stopwords. In the current implementation, we use word-based uni-gram indexing for both English and Japanese documents. In other words, compound words are decomposed into base words in the surrogates. Note that indexing and retrieval methods are theoretically independent of

the translation method.

Thereafter, the “translator” processes a query in the source language (“S-query”) to output the translation (“T-query”). T-query can consist of more than one translation, because multiple translations are often appropriate for a single technical term.

Finally, the “IR engine” computes the similarity between T-query and each document in the surrogates based on the vector space model (Salton and McGill, 1983), and sorts document according to the similarity, in descending order. We compute term weight based on the notion of TF-IDF. Note that T-query is decomposed into base words, as performed in the document preprocessing.

In Section 3, we will explain the “translator” in Figure 1, which involves compound word translation and transliteration modules.

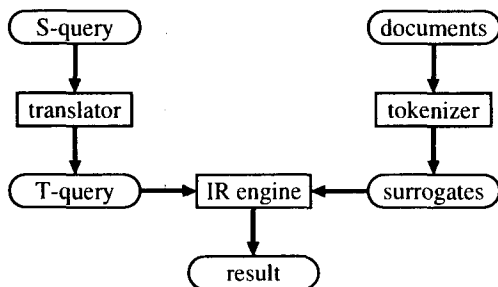


Figure 1: The overall design of our CLIR system

### 3 Translation Module

#### 3.1 Overview

Given a query in the source language, tokenization is first performed as for target documents (see Figure 1). To put it more precisely, we use WordNet and ChaSen for English and Japanese queries, respectively. We then discard stopwords and extract only content words. Here, “content words” refer to both single and compound words. Let us take the following query as an example:

improvement of data mining methods.

For this query, we discard “of”, to extract “improvement” and “data mining methods”.

Thereafter, we translate each extracted content word individually. Note that we currently do not consider relation (e.g. syntactic relation and collocational information) between content words. If a single word, such as “improvement” in the example above, is listed in our bilingual dictionary (we will explain the way to produce the dictionary in Section 3.2), we use all possible translation candidates as query terms for the subsequent retrieval phase.

Otherwise, compound word translation is performed. In the case of Japanese-English translation, we consider all possible segmentations of the input word, by consulting the dictionary. Then, we select such segmentations that consist of the minimal number of base words. During the segmentation process, the dictionary derives all possible translations for base words. At the same time, transliteration is performed whenever *katakana* sequences unlisted in the dictionary are found. On the other hand, in the case of English-Japanese translation, transliteration is applied to any unlisted base word (including the case where the input English word consists of a single base word). Finally, we compute the probability of occurrence of each combination of base words in the target language, and select those with greater probabilities, for both Japanese-English and English-Japanese translations.

#### 3.2 Compound Word Translation

This section briefly explains the compound word translation method we previously proposed (Fujii and Ishikawa, 1999). This method translates input compound words on a word-by-word basis, maintaining the word order in the source language<sup>2</sup>. The formula for the source compound word and one translation candidate are represented as below.

$$\begin{aligned}
 S &= s_1, s_2, \dots, s_n \\
 T &= t_1, t_2, \dots, t_n
 \end{aligned}$$

<sup>2</sup>A preliminary study showed that approximately 95% of compound technical terms defined in a bilingual dictionary maintain the same word order in both source and target languages.

Here,  $s_i$  and  $t_i$  denote  $i$ -th base words in source and target languages, respectively. Our task, i.e., to select  $T$  which maximizes  $P(T|S)$ , is transformed into Equation (1) through use of the Bayesian theorem.

$$\arg \max_T P(T|S) = \arg \max_T P(S|T) \cdot P(T) \quad (1)$$

$P(S|T)$  and  $P(T)$  are approximated as in Equation (2), which has commonly been used in the recent statistical NLP research (Church and Mercer, 1993).

$$\begin{aligned} P(S|T) &\approx \prod_{i=1}^n P(s_i|t_i) \\ P(T) &\approx \prod_{i=1}^{n-1} P(t_{i+1}|t_i) \end{aligned} \quad (2)$$

We produced our own dictionary, because conventional dictionaries are comprised primarily of general words and verbose definitions aimed at human readers. We extracted 59,533 English/Japanese translations consisting of *two* base words from the EDR technical terminology dictionary, which contains about 120,000 translations related to the information processing field (Japan Electronic Dictionary Research Institute, 1995), and segment Japanese entries into two parts<sup>3</sup>. For this purpose, simple heuristic rules based mainly on Japanese character types (i.e., *kanji*, *katakana*, *hiragana*, alphabets and other characters like numerals) were used. Given the set of compound words where Japanese entries are segmented, we correspond English-Japanese base words on a word-by-word basis, maintaining the word order between English and Japanese, to produce a Japanese-English/English-Japanese base word dictionary. As a result, we extracted 24,439 Japanese base words and 7,910 English base words from the EDR dictionary. During the dictionary production, we also count the collocational frequency for each combination of  $s_i$  and  $t_i$ , in order to estimate  $P(s_i|t_i)$ . Note that in the case where

<sup>3</sup>The number of base words can easily be identified based on English words, while Japanese compound words lack lexical segmentation.

$s_i$  is *transliterated* into  $t_i$ , we use an arbitrarily predefined value for  $P(s_i|t_i)$ . For the estimation of  $P(t_{i+1}|t_i)$ , we use the word-based bi-gram statistics obtained from target language corpora, i.e., “documents” in the collection (see Figure 1).

### 3.3 Transliteration

Figure 2 shows example correspondences between English and (romanized) *katakana* words, where we insert hyphens between each *katakana* character for enhanced readability. The basis of our transliteration method is analogous to that for compound word translation described in Section 3.2. The formula for the source word and one transliteration candidate are represented as below.

$$\begin{aligned} S &= s_1, s_2, \dots, s_n \\ T &= t_1, t_2, \dots, t_n \end{aligned}$$

However, unlike the case of compound word translation,  $s_i$  and  $t_i$  denote  $i$ -th “symbols” (which consist of one or more letters), respectively. Note that we consider only such  $T$ ’s that are indexed in the inverted file, because our transliteration method often outputs a number of incorrect words with great probabilities. Then, we compute  $P(T|S)$  for each  $T$  using Equations (1) and (2) (see Section 3.2), and select  $k$ -best candidates with greater probabilities. The crucial content here is the way to produce a bilingual dictionary for *symbols*. For this purpose, we used approximately 3,000 *katakana* entries and their English translations listed in our base word dictionary. To illustrate our dictionary production method, we consider Figure 2 again. Looking at this figure, one may notice that the first letter in each *katakana* character tends to be contained in its corresponding English word. However, there are a few exceptions. A typical case is that since Japanese has no distinction between “L” and “R” sounds, the two English sounds collapse into the same Japanese sound. In addition, a single English letter corresponds to multiple *katakana* characters, such as “x” to “*ki-su*” in “<text, *te-ki-su-to*>”. To sum up, English and

romanized *katakana* words are not exactly identical, but *similar* to each other.

English	<i>katakana</i>
system	<i>shi-su-te-mu</i>
mining	<i>ma-i-ni-n-gu</i>
data	<i>dee-ta</i>
network	<i>ne-tto-waa-ku</i>
text	<i>te-ki-su-to</i>
collocation	<i>ko-ro-ke-i-sho-n</i>

Figure 2: Examples of English-*katakana* correspondence

We first manually define the similarity between the English letter *e* and the first romanized letter for each *katakana* character *j*, as shown in Table 1. In this table, “phonetically similar” letters refer to a certain pair of letters, such as “L” and “R”<sup>4</sup>. We then consider the similarity for any possible combination of letters in English and romanized *katakana* words, which can be represented as a matrix, as shown in Figure 3. This figure shows the similarity between letters in “<text, *te-ki-su-to*>”. We put a dummy letter “\$”, which has a positive similarity only to itself, at the end of both English and *katakana* words. One may notice that matching plausible symbols can be seen as finding the path which maximizes the total similarity from the first to last letters. The best path can easily be found by, for example, Dijkstra’s algorithm (Dijkstra, 1959). From Figure 3, we can derive the following correspondences: “<te, *te*>”, “<x, *ki-su*>” and “<t, *to*>”. The resultant correspondences contain 944 Japanese and 790 English symbol types, from which we also estimated  $P(s_i|t_i)$  and  $P(t_{i+1}|t_i)$ .

As can be predicted, a preliminary experiment showed that our transliteration method is not accurate when compared with a word-based translation. For example, the Japanese word “*re-ji-su-ta* (register)” is transliterated to “resister”, “resistor” and “register”, with the probability score in descending order. How-

<sup>4</sup>We identified approximately twenty pairs of phonetically similar letters.

ever, combined with the compound word translation, irrelevant transliteration outputs are expected to be discarded. For example, a compound word like “*re-ji-su-ta tensou gengo* (register transfer language)” is successfully translated, given a set of base words “*tensou* (transfer)” and “*gengo* (language)” as a context.

Table 1: The similarity between English and Japanese letters

condition	similarity
<i>e</i> and <i>j</i> are identical	3
<i>e</i> and <i>j</i> are phonetically similar	2
both <i>e</i> and <i>j</i> are vowels or consonants	1
otherwise	0




Figure 3: An example matrix for English-Japanese symbol matching (arrows denote the best path)

## 4 Evaluation

This section investigates the performance of our CLIR system based on the TREC-type evaluation methodology: the system outputs 1,000 top documents, and TREC evaluation software is used to calculate the recall-precision trade-off and 11-point average precision.

For the purpose of our evaluation, we used the NACSIS test collection (Kando et al., 1998). This collection consists of 21 Japanese queries and approximately 330,000 documents (in ci-

ther a combination of English and Japanese or either of the languages individually), collected from technical papers published by 65 Japanese associations for various fields. Each document consists of the document ID, title, name(s) of author(s), name/date of conference, hosting organization, abstract and keywords, from which titles, abstracts and keywords were used for our evaluation. We used as target documents approximately 187,000 entries where abstracts are in both English and Japanese. Each query consists of the title of the topic, description, narrative and list of synonyms, from which we used only the description. Roughly speaking, most topics are related to electronic, information and control engineering. Figure 4 shows example descriptions (translated into English by one of the authors). Relevance assessment was performed based on one of the three ranks of relevance, i.e., “relevant”, “partially relevant” and “irrelevant”. In our evaluation, relevant documents refer to both “relevant” and “partially relevant” documents<sup>5</sup>.

ID	description
0005	dimension reduction for clustering
0006	intelligent information retrieval
0019	syntactic analysis methods for Japanese
0024	machine translation systems

Figure 4: Example descriptions in the NACSIS query

#### 4.1 Evaluation of compound word translation

We compared the following query translation methods:

- (1) a control, in which all possible translations derived from the (original) EDR technical terminology dictionary are used as query terms (“EDR”),
- (2) all possible base word translations derived from *our* dictionary are used (“all”),
- (3) randomly selected  $k$  translations derived from our bilingual dictionary are used (“random”),
- (4)  $k$ -best translations through compound word translation are used (“CWT”).

(3) randomly selected  $k$  translations derived from our bilingual dictionary are used (“random”),

(4)  $k$ -best translations through compound word translation are used (“CWT”).

For system “EDR”, compound words unlisted in the EDR dictionary were manually segmented so that substrings (shorter compound words or base words) can be translated. For both systems “random” and “CWT”, we arbitrarily set  $k = 3$ . Figure 5 and Table 2 show the recall-precision curve and 11-point average precision for each method, respectively. In these, “J-J” refers to the result obtained by the Japanese-Japanese IR system, which uses as documents Japanese titles/abstracts/keywords comparable to English fields in the NACSIS collection. This can be seen as the upper bound for CLIR performance<sup>6</sup>. Looking at these results, we can conclude that the dictionary production and probabilistic translation methods we proposed are effective for CLIR.




Figure 5: Recall-Precision curves for evaluation of compound word translation

<sup>5</sup>The result did not significantly change depending on whether we regarded “partially relevant” as relevant or not.

<sup>6</sup>Regrettably, since the NACSIS collection does not contain English queries, we cannot estimate the upper bound performance by English-English IR.

Table 2: Comparison of average precision for evaluation of compound word translation

	avg. precision	ratio to J-J
J-J	0.204	—
CWT	0.193	0.946
all	0.171	0.838
EDR	0.130	0.637
random	0.116	0.569

## 4.2 Evaluation of transliteration

In the NACSIS collection, three queries contain *katakana* (base) words unlisted in our bilingual dictionary. Those words are “*ma-i-ni-n-gu* (mining)” and “*ko-ro-ke-i-sho-n* (collocation)”. However, to emphasize the effectiveness of transliteration, we compared the following extreme cases:

- (1) a control, in which every *katakana* word is discarded from queries (“control”),
- (2) a case where transliteration is applied to every *katakana* word and top 10 candidates are used (“translit”).

Both cases use system “CWT” in Section 4.1. In the case of “translit”, we do not use *katakana* entries listed in the base word dictionary. Figure 6 and Table 3 show the recall-precision curve and 11-point average precision for each case, respectively. In these, results for “CWT” correspond to those in Figure 5 and Table 2, respectively. We can conclude that our transliteration method significantly improves the baseline performance (i.e., “control”), and comparable to word-based translation in terms of CLIR performance.

An interesting observation is that the use of transliteration is robust against *typos* in documents, because a number of similar strings are used as query terms. For example, our transliteration method produced the following strings for “*ri-da-ku-sho-n* (reduction)”:

riduction, redction, redaction, reduction.

All of these words are effective for retrieval, because they are contained in the target documents.




Figure 6: Recall-Precision curves for evaluation of transliteration

Table 3: Comparison of average precision for evaluation of transliteration

	avg. precision	ratio to J-J
J-J	0.204	—
CWT	0.193	0.946
translit	0.193	0.946
control	0.115	0.564

## 4.3 Evaluation of the overall performance

We compared our system (“CWT+translit”) with the Japanese-Japanese IR system, where (unlike the evaluation in Section 4.2) transliteration was applied only to “*ma-i-ni-n-gu* (mining)” and “*ko-ro-ke-i-sho-n* (collocation)”. Figure 7 and Table 4 show the recall-precision curve and 11-point average precision for each system, respectively, from which one can see that our CLIR system is quite comparable with the monolingual IR system in performance. In addition, from Figure 5 to 7, one can see that the monolingual system generally performs better

at lower recall while the CLIR system performs better at higher recall.

For further investigation, let us discuss similar experimental results reported by Kando and Aizawa (1998), where a bilingual dictionary produced from Japanese/English keyword pairs in the NACSIS documents is used for query translation. Their evaluation method is almost the same as performed in our experiments. One difference is that they use the “OpenText” search engine<sup>7</sup>, and thus the performance for Japanese-Japanese IR is higher than obtained in our evaluation. However, the performance of their Japanese-English CLIR systems, which is roughly 50-60% of that for *their* Japanese-Japanese IR system, is comparable with our CLIR system performance. It is expected that using a more sophisticated search engine, our CLIR system will achieve a higher performance than that obtained by Kando and Aizawa.




Figure 7: Recall-Precision curves for evaluation of overall performance

## 5 Conclusion

In this paper, we proposed a Japanese/English cross-language information retrieval system, targeting technical documents. We combined a query translation module, which performs

<sup>7</sup>Developed by OpenText Corp.

Table 4: Comparison of average precision for evaluation of overall performance

	avg. precision	ratio to J-J
J-J	0.204	--
CWT + translit	0.212	1.04

compound word translation and transliteration, with an existing monolingual retrieval method. Our experimental results showed that compound word translation and transliteration methods individually improve on the baseline performance, and when used together the improvement is even greater. Future work will include the application of automatic word alignment methods (Fung, 1995; Smadja et al., 1996) to enhance the dictionary.

## Acknowledgments

The authors would like to thank Noriko Kando (National Center for Science Information Systems, Japan) for her support with the NACSIS collection.

## References

- Lisa Ballesteros and W. Bruce Croft. 1998. Resolving ambiguity for cross-language retrieval. In *Proceedings of the 21th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 64–71.
- Jaime G. Carbonell, Yiming Yang, Robert E. Frederking, Ralf D. Brown, Yibing Geng, and Danny Lee. 1997. Translingual information retrieval: A comparative evaluation. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, pages 708–714.
- Hsin-Hsi Chen, Sheng-Jie Huang, Yung-Wei Ding, and Shih-Chung Tsai. 1998. Proper name translation in cross-language information retrieval. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and the 17th International Conference on Computational Linguistics*, pages 232–236.
- Kenneth W. Church and Robert L. Mercer. 1993. Introduction to the special issue on computational linguistics using large corpora. *Computational Linguistics*, 19(1):1–24.
- Mark W. Davis and William C. Ogden. 1997. QUILT: Implementing a large-scale cross-language



- text retrieval system. In *Proceedings of the 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 92–98.
- Edsger W. Dijkstra. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271.
- Susan T. Dumais, Thomas K. Landauer, and Michael L. Littman. 1996. Automatic cross-linguistic information retrieval using latent semantic indexing. In *ACM SIGIR Workshop on Cross-Linguistic Information Retrieval*.
- Atsushi Fujii and Tetsuya Ishikawa. 1999. Cross-language information retrieval using compound word translation. In *Proceedings of the 18th International Conference on Computer Processing of Oriental Languages*, pages 105–110.
- Pascale Fung. 1995. A pattern matching method for finding noun and proper noun translations from noisy parallel corpora. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, pages 236–243.
- Denis A. Gachot, Elke Lange, and Jin Yang. 1996. The SYSTRAN NLP browser: An application of machine translation technology in multilingual information retrieval. In *ACM SIGIR Workshop on Cross-Linguistic Information Retrieval*.
- David A. Hull and Gregory Grefenstette. 1996. Querying across languages: A dictionary-based approach to multilingual information retrieval. In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 49–57.
- Japan Electronic Dictionary Research Institute. 1995. Technical terminology dictionary (information processing). (In Japanese).
- Noriko Kando and Akiko Aizawa. 1998. Cross-lingual information retrieval using automatically generated multilingual keyword clusters. In *Proceedings of the 3rd International Workshop on Information Retrieval with Asian Languages*, pages 86–94.
- Noriko Kando, Teruo Koyama, Keizo Oyama, Kyo Kageura, Masaharu Yoshioka, Toshihiko Nozue, Atsushi Matsumura, and Kazuko Kuriyama. 1998. NTCIR: NACSIS test collection project. In *The 20th Annual BCS-IRSG Colloquium on Information Retrieval Research*.
- Kevin Knight and Jonathan Graehl. 1998. Machine transliteration. *Computational Linguistics*, 24(4):599–612.
- Yuji Matsumoto, Akira Kitauchi, Tatsuo Yamashita, Osamu Imaichi, and Tomoaki Imamura. 1997. Japanese morphological analysis system ChaSen manual. Technical Report NAIST-IS-TR97007, NAIST. (In Japanese).
- George A. Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, Katherine Miller, and Randee Teng. 1993. Five papers on WordNet. Technical Report CLS-Rep-43, Cognitive Science Laboratory, Princeton University.
- P. E. Mongar. 1969. International co-operation in abstracting services for road engineering. *The Information Scientist*, 3:51–62.
- Douglas W. Oard and Paul Hackett. 1997. Document translation for cross-language text retrieval at the University of Maryland. In *The 6th Text Retrieval Conference*.
- Akitoshi Okumura, Kai Ishikawa, and Kenji Satoh. 1998. Translingual information retrieval by a bilingual dictionary and comparable corpus. In *The 1st International Conference on Language Resources and Evaluation, workshop on translingual information management: current levels and future abilities*.
- Ari Pirkola. 1998. The effects of query structure and dictionary setups in dictionary-based cross-language information retrieval. In *Proceedings of the 21th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 55–63.
- Gerard Salton and Michael J. McGill. 1983. *Introduction to Modern Information Retrieval*. McGraw-Hill.
- Gerard Salton. 1970. Automatic processing of foreign language documents. *Journal of the American Society for Information Science*, 21(3):187–194.
- Páraic Sheridan and Jean Paul Ballerini. 1996. Experiments in multilingual information retrieval using the SPIDER system. In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 58–65.
- Frank Smadja, Kathleen R. McKeown, and Vasileios Hatzivassiloglou. 1996. Translating collocations for bilingual lexicons: A statistical approach. *Computational Linguistics*, 22(1):1–38.
- Stephen Wan and Cornelia Maria Verspoor. 1998. Automatic English-Chinese name transliteration for development of multilingual resources. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and the 17th International Conference on Computational Linguistics*, pages 1352–1356.